

When you're evaluating an AI workspace, a limit such as "100 files each" can sound clear—until you try to upload a spreadsheet, a folder of research, or a collection of PDFs and wonder what the product counts. Is a file counted once per project, once per conversation, or across your whole account? Does a ZIP archive count as one file or as the documents inside it? And does "15 MB files" refer to the upload limit for each item or the total size of everything you attach?

Those questions matter if you use Suprmind for investment research, product planning, policy review, or other decision-heavy work. The practical answer depends on the exact wording and scope of the limit shown in your account. Treat "100 files each" and "15 MB files" as separate constraints, then confirm where each applies before designing a workflow around them.

This guide explains how to interpret those limits, what generally counts as a file, how to plan for large research collections, and why limits are only one part of comparing Suprmind's plan-based approach with AI services that bill by usage.

Start with the two limits

"100 files each" describes a file-count ceiling, while "15 MB files" describes a per-file size ceiling. Neither phrase, by itself, tells you whether the count is scoped to a project, a thread, a workspace, or an account. That scope is the first thing to verify in Suprmind's current plan details or in the upload interface.

- **100 files each:** A maximum number of file items in the scope named by the product. Check whether that means each project, thread, or another container.
- **15 MB files:** A maximum size for an individual upload, if the limit is stated per file. Confirm whether the product uses MB or MiB and whether it measures the original file or an uploaded version.
- **Combined uploads:** A separate total-size limit may apply, even when every individual file is below 15 MB.

For example, 100 files that are each 1 MB would meet a 15 MB per-file limit, but they could still exceed a separate total-storage or batch-upload limit. Conversely, a single 16 MB PDF could be rejected even if you have only uploaded one file. Per-file, per-project, and total-account limits are different kinds of rules.

What counts as a file?

In most upload systems, "a file" means an individual item you select or drag into the uploader—not a page, paragraph, worksheet, or research source in the abstract. A PDF with 200 pages is typically one uploaded file. A spreadsheet with several tabs is typically one file. Ten separate PDFs are typically ten files.


That is a useful starting point, but product behavior can vary. Suprmind may process or index a file in multiple internal steps without counting those steps as extra user uploads. The reliable unit is usually the item shown in the project or thread's attachment list, but check the interface and plan documentation rather than assuming that internal processing maps one-to-one to the displayed quota.

Common examples

- **One PDF:** Usually one file, regardless of how many pages it contains.
- **One XLSX workbook:** Usually one file, even if it contains multiple sheets.
- **Several separate images:** Usually one file per image.

- **A folder:** A folder may not be a supported upload unit. If you select its contents, the individual items may count separately.
- **A ZIP archive:** It may count as one uploaded archive, but its contents may not be readable, supported, or counted the same way. Do not assume that compressing 200 documents lets you bypass a 100-file limit.
- **A revised document:** Uploading a new version may create another file item, depending on whether the product replaces the previous version or adds a new attachment.
- **A link to a document:** A URL is not necessarily an uploaded file. A connected-source or web-import feature may have different rules.

For a decision workflow, think in terms of the actual items you need the AI to read. If you have 80 reports plus 25 supporting exhibits, you may have 105 files—not one “research package.” If you combine relevant materials into a single supported document, you might reduce the attachment count, but you should preserve source labels and logical sections so the resulting analysis remains auditable.



How to plan around the 15 MB limit

A 15 MB per-file limit can be plenty for text-heavy documents and restrictive for image-heavy material. File size is not a reliable proxy for usefulness: a scanned PDF with high-resolution pages can be large while containing little searchable text, whereas a much smaller text PDF may be easy to analyze.

Before uploading, remove irrelevant pages, duplicate copies, and unused exhibits. If a file is too large, consider exporting only the needed sections or reducing image resolution where that does not impair legibility. Avoid splitting one source into many tiny fragments without a reason; excessive fragmentation can make it harder to preserve context and citations.

1. **Inventory the source set.** List the documents, their formats, and their sizes before starting a project.
2. **Prioritize evidence.** Upload primary sources and decision-critical documents first; keep background material separate.
3. **Check text quality.** For scanned documents, verify that text is searchable or that the product supports the relevant image or OCR workflow.
4. **Test a representative sample.** Use one ordinary document and one edge case—such as a large spreadsheet or scanned PDF—to confirm the upload and extraction behavior.
5. **Record provenance.** Keep original filenames and dates where possible, and note any document that was excluded, compressed, or split.

Do “100 files each” and unlimited projects mean unlimited research?

No. “Unlimited projects” generally refers to the ability to create many project containers; it does not automatically mean unlimited files, storage, processing, or model usage. A product can allow unlimited projects while applying a separate per-project attachment cap, a per-file size ceiling, or other usage limits. Confirm how Suprmind defines “each” and whether the file allowance resets for every project or is shared across the account.

This distinction affects how you organize work. If the cap is per project, one project for each investment target or policy review may be practical. If the cap is account-wide, dividing a corpus across many projects may not increase your total allowance. And even if every project has its own allowance, scattering a single decision record across multiple projects could make comparison and traceability harder.

Before relying on unlimited projects as a way to organize large research programs, check:

- Whether project creation is unlimited on your specific plan.
- Whether the 100-file allowance applies per project, per thread, or per account.
- Whether files can be reused across projects without making duplicate uploads.
- Whether an upload limit is separate from storage, indexing, or monthly AI usage.
- Whether deleted or replaced files continue to count toward a limit.

Why multi-model orchestration changes the value calculation

A shared thread that can orchestrate multiple models is valuable for more than convenience. It can let a team frame one question, provide common evidence, compare responses from models such as OpenAI and Anthropic, and then use structured validation to examine where their conclusions agree or differ. The key benefit is a more consistent decision workflow—not simply access to a longer list of model names.

For example, a strategy team might ask two models to identify assumptions in an investment memo, then compare the results against a structured checklist: market size, customer evidence, unit economics, competitive risks, and missing data. A product team could use the same pattern to evaluate competing roadmap proposals. A compliance team might ask for clauses that require review, then validate those findings against a defined policy rubric.

Shared context helps reduce the overhead of copying prompts and documents between separate tools. Structured validation helps turn fluent answers into reviewable work: a team can specify required fields, ask for evidence, flag uncertainty, and check that the output follows a consistent format. These steps do not eliminate the need for human review. They make review easier to repeat and document.

Use a validation checklist

- Does each conclusion point to a relevant source or clearly state that support is missing?
- Are facts separated from assumptions, interpretations, and recommendations?
- Did the models identify conflicting evidence or unresolved questions?
- Are the required output fields present and in the requested structure?
- Has a subject-matter reviewer checked high-impact claims?

That is where a file cap becomes a workflow consideration. If a project can hold only a finite number of files, prioritize the evidence needed for the decision and maintain a separate source inventory. More files do not automatically produce better analysis; a focused, well-labeled evidence set often works better than a disorganized dump.

Plan-based pricing versus per-token billing

Suprmind's plan-based pricing and a provider's per-token billing answer different purchasing questions. A plan price can make budgeting more predictable, while a usage-based API bill can vary with the number of requests, model choice, input length, and generated output. But a subscription price does not, on its own, tell you how many files you can upload or which features are included. Those are determined [what is multi-model ai](#) by plan terms and feature gates.

When comparing tools, evaluate the complete operating cost: subscription, usage, seats, integrations, administration, and the time required to assemble and review outputs. A lower monthly price may be attractive for an individual test, while a higher plan may be justified if it includes collaboration, model orchestration, structured workflows, or controls your team actually needs.

As a simple price reference, **Spark - \$19/month** is an example of a plan price. Treat that as a pricing example, not proof that every capability or limit discussed here is included at that tier. Verify the current feature list and quota for Spark before you commit.

Feature gating: Spark, Pro, Frontier, Power, and Enterprise

Plan names often signal that capabilities are divided into tiers, but names alone do not establish what each tier includes. Spark, Pro, Frontier, Power, and Enterprise may differ in access to models, orchestration, collaboration, administrative controls, or usage allowances. The relevant question is not simply "Which plan is best?" It is "Which plan supports the workflow we need, at the limits we can operate within?"

Evaluation area What to verify across Spark, Pro, Frontier, Power, and Enterprise File limits Whether 100 files each and 15 MB files apply, and exactly what “each” refers to. Model access Which models are available and whether model choice varies by plan. Orchestration Whether multiple models can work within a shared thread or workflow. Decision workflows Whether structured validation, reusable templates, or other review features are included. Team controls Whether collaboration, permissions, administrative features, and support meet your requirements. Usage and cost Whether the plan has included usage, separate metering, or other limits that affect total cost.

For a small team, a lower tier may be sufficient for testing a focused workflow. A larger research or compliance team may care more about shared workspaces, consistent review, administration, and predictable limits. Enterprise procurement should also ask about security, data handling, retention, support, and contractual terms rather than focusing on file counts alone.

How to compare adjacent business tools

AI research workspaces are not substitutes for every business system. When comparing software budgets, separate the core job each tool performs. For example, **Gaviti** is AR software for accounts receivable workflows; its pricing is by quote, with a free trial and demo available. **Paylocity** is a separate business platform with custom pricing and no free trial. **Zoho Books** is accounting software listed at \$15 per month, with a 14-day free trial. These examples can help anchor a broader software evaluation, but their pricing and trial terms should be verified for your region and current offer.

The point is not to compare their feature sets directly with Suprmind. It is to compare buying models: quote-based sales, custom pricing, and published subscription pricing each create different evaluation paths. For any vendor, confirm what is included, what requires a higher tier, and what happens when you reach a usage limit.

A practical pre-purchase test

Use a small but realistic workflow to test the limits before moving a live decision process into the product. Choose a project with a known source set and a clear deliverable—for example, a short investment memo, a product roadmap comparison, or a policy clause review.

1. Count the source files and note the size of each one.
2. Upload a representative set, including any formats your team relies on.
3. Observe whether the interface identifies a project, thread, or account-level limit.
4. Run a shared-thread task using the model options available on your plan.
5. Ask for a structured output with evidence, uncertainty, and missing information.
6. Check whether the result can be validated and reviewed by another team member.
7. Ask the vendor what happens when a file is too large or the count limit is reached.

Bottom line

Interpret “100 files each” as a file-count limit whose scope still needs confirmation, and “15 MB files” as a separate size constraint whose exact measurement and exceptions should be checked. A document’s pages or spreadsheet tabs usually do not become separate files, but folders, archives, duplicate versions, and linked sources can behave differently. Most importantly, unlimited projects does not mean unlimited uploads or usage.

For Suprmind, evaluate the limits alongside multi-model orchestration, shared-thread context, structured validation, and plan-based feature access. Confirm how Spark, Pro, Frontier, Power, and Enterprise differ before

choosing a tier. Then test your real evidence set—not an idealized demo—so you know whether the file rules support the decisions your team needs to make.